

Performance and Scalability Evaluation of “Big Memory” on Blue Gene Linux *

Kazutomo Yoshii¹, Kamil Iskra¹, Harish Naik¹, Pete Beckman^{1,2}, P. Chris Broekema³

¹Mathematics and Computer Science Division

²Leadership Computing Facility

Argonne National Laboratory

9700 South Cass Avenue, Argonne, IL 60439, USA

³ASTRON

Netherlands Institute for Radio Astronomy

Oude Hoogeveensedijk 4, Dwingeloo, The Netherlands

*This work was supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

LOFAR is funded by the Dutch government in the BSIK programme for interdisciplinary research for improvements of the knowledge infrastructure. Additional funding is provided by the European Union, European Regional Development Fund (EFRO) and by the “Samenwerkingsverband Noord-Nederland,” EZ/KOMPAS.

Proposed running head: **Big Memory on Blue Gene Linux**

Kazutomo Yoshii
Mathematics and Computer Science Division
Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439, USA
phone: +1-630-252-4476
fax: +1-630-252-5986
email: kazutomo@mcs.anl.gov
(corresponding author)

Kamil Iskra
Mathematics and Computer Science Division
Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439, USA
email: iskra@mcs.anl.gov
phone: +1-630-252-7197
fax: +1-630-252-5986

Harish Naik
Mathematics and Computer Science Division
Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439, USA
email: hnaik@mcs.anl.gov
phone: +1-630-252-2029
fax: +1-630-252-5986

Pete Beckman
Leadership Computing Facility and
Mathematics and Computer Science Division
Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439, USA
email: beckman@mcs.anl.gov
phone: +1-630-252-9020
fax: +1-630-252-5986

P. Chris Broekema
Netherlands Institute for Radio Astronomy
Oude Hoogeveensedijk 4,
Dwingeloo, The Netherlands
email: broekema@astron.nl

Abstract

We address memory performance issues observed in Blue Gene Linux, and discuss the design and implementation of “Big Memory”—an alternative, transparent memory space introduced to eliminate the memory performance issues. We evaluate the performance of Big Memory using custom memory benchmarks, NAS Parallel Benchmarks, and the Parallel Ocean Program, at a scale of up to 4096 nodes. We find that Big Memory is successful in resolving performance issues we had encountered. For the ocean simulation program, we even find that Linux provides better scalability than does the lightweight compute node kernel designed solely for high-performance applications. Originally intended exclusively for compute node tasks, our new memory subsystem turns out to dramatically improve the performance of certain I/O node applications as well. We demonstrate this performance using the central processor of the LOw Frequency ARray radio telescope as an example.

Keywords: Linux, Blue Gene, OS kernel, memory performance, TLB

1 Introduction

The Blue Gene architecture [9, 14], developed by IBM, is one of the most successful contemporary massively parallel computer architectures, thanks to a high-speed interconnect, a highly scalable design, and a very low power consumption compared to other supercomputers.

Blue Gene machines normally run a dedicated compute node kernel (CNK) [15] on the compute nodes. CNK is essentially a microkernel that supports only one user thread per CPU core and provides a simplified, offset-based mapping from physical memory to the virtual address space. This design keeps the kernel small and simple; more important, since Blue Gene lacks hardware to handle translation lookaside buffer (TLB) misses efficiently, it maximizes the memory access performance, as well as the floating-point performance. It also simplifies the programming of hardware devices, in particular the DMA engine discussed later.

Unfortunately, the simplicity of the design is also an obstacle; it brings an inflexibility and a lack of features that are generally taken for granted in more general-purpose operating system kernels, such as multitasking and time sharing. This situation prompted us to replace CNK with a Linux kernel as a part of the ZeptoOS project [19], in an effort to create a fully open software stack to enable independent computer science research on massively parallel architectures, enhance community collaboration, and foster innovation.

In previous publications related to ZeptoOS, we focused on operating system jitter [2, 4] and on I/O forwarding [11]. This paper focuses on the performance of our Linux-based ZeptoOS compute node kernel, with an emphasis on memory management.

In the remainder of this section, we outline the key characteristics of the Blue Gene hardware design and their consequences for the software stack (both CNK and the ZeptoOS one). We focus on the Blue Gene/P (BG/P) architecture [10] introduced in 2007 to replace the original Blue Gene/L design [14].

Blue Gene racks consist of two kinds of nodes: compute nodes, running the application code, and I/O nodes, responsible for system services such as file I/O. These two types of nodes feature the same hardware; however, the available interconnects are different. The primary interconnect available on the compute nodes is a 3D torus network, used for high-performance point-to-point communication. A collective network is also available for more sophisticated operations; it doubles as a file I/O backbone connecting compute nodes with I/O nodes. Unlike the compute nodes, which normally run CNK, I/O nodes run a Linux kernel. I/O nodes also have a 10 GbE link to connect them to file servers, login nodes, and the service node. The only practical means of communication between compute nodes and the outside world is through the I/O nodes, using I/O

forwarding.

We chose Linux as the foundation of our ZeptoOS kernel for the compute nodes because of its openness and popularity, but our decision was made easier by the fact that Linux already ran on the nearly identical I/O nodes. Only minor changes to the I/O node kernel itself were needed to boot it on the compute nodes. However, a number of far-reaching changes were required to make it *useful*. For example, we had to replace the CNK-specific IBM I/O forwarding and job control software with ZOID [11]. ZOID enables access to remote file systems from the compute nodes running Linux, but it also provides an interactive login capability to the compute nodes from the outside, which proved invaluable in further software development.

Internally, BG/P nodes use PowerPC 450 CPUs—a quad-core, 32-bit design with SMP support, running at 850 MHz. Each processor core has a dual-pipeline floating-point unit with *fused multiply-add* (FMA) instructions. The peak floating-point performance of the whole CPU is 13.6 Gflops. Each core has a 32 KB¹ L1 instruction cache and a 32 KB L1 data cache (the latter featuring a snoop filter to provide cache coherency between cores). The peak fill rate is 6.8 GB/s. The L2 cache is smaller than the L1 and serves as a stream prefetching buffer. The CPU has a common 8 MB L3 cache. Nodes have either 2 GB or 4 GB of main memory. The main store bandwidth is 13.6 GB/s.

At the time of writing, PowerPC 450 is not a component available separately on the market; it can be purchased only as part of a Blue Gene/P system. Consequently, a stock Linux kernel does not have the support needed to make Linux boot on that CPU. However, starting with Blue Gene/P, IBM provides the necessary patches to enable Linux to work on this processor. Regular 32-bit PowerPC executables run well on a BG/P Linux kernel; however, executables compiled specifically for BG/P using a patched GNU C compiler might not work on other 32-bit PowerPC processors because of the custom BG/P FMA instruction set.

Because of its embedded systems origins, the processor has only 64 TLB entries per core. Even worse, TLB misses must be handled in software, inside the Linux kernel, with an average cost of approximately 0.3 μ s. Hence, with the default PowerPC Linux page size of 4 KB, if the memory is accessed randomly, the working set of the program must not be larger than 256 KB before the performance significantly degrades. The processor itself supports various page sizes ranging from 1 KB to 1 GB. Pages of different sizes can be used simultaneously; unfortunately, the Linux kernel lacks the flexibility needed to take advantage of this feature. We have experimented with increasing the system page size to 64 KB; we show the results later in

¹Throughout this paper, we use KB, MB, or GB in the context of memory size; 1 KB equals 1,024 bytes.

the paper. Unfortunately, this approach is not really an option on the I/O nodes because the legacy software running there (in particular, the GPFS file system client code) works only with 4 KB pages. CNK, on the other hand, takes full advantage of the hardware and statically maps all the system memory using large TLB entries, thereby eliminating TLB misses.

Unlike its predecessor used in Blue Gene/L, the PowerPC 450 has a coherent L1 data cache; the PowerPC *load and reserve* instruction works between cores, and coherent pages are available. However, the `tlbsync` instruction is not supported, so the Linux kernel has to use an interprocess-interrupt (IPI) to synchronize software TLB management. The BG/P Linux kernel is configured for a virtual memory split of 3 GB user space and 1 GB kernel space, with HIGHMEM enabled by default. BG/P-specific additions primarily include various network drivers.

In Blue Gene/P, the torus network between the compute nodes has been enhanced with a DMA engine. The engine can deal only with physical addresses; a software layer has to translate virtual addresses to physical ones. This process is simple with IBM's CNK, where contiguous virtual addresses are also contiguous in physical space. Unfortunately, with paged memory used in Linux, this is not the case—the translation is a lot more complex, and fragmented physical address space seriously limits the size of DMA operations, hurting performance.

Given the problems caused by paged memory on BG/P, we decided to investigate an alternative memory management scheme.

2 Related Work

Linux does provide support for large memory pages, through the *hugetlbfs* [7] mechanism. Using these pages dramatically reduces the number of TLB misses, improving performance. However, this feature is not transparent—applications need to invoke the `mmap` system call explicitly to make that memory available, and by then it is too late to use the memory for segments such as application text, heap, or stack.

Shmueli et al. [18] evaluated Linux on the compute nodes of Blue Gene/L and identified TLB misses as a major source of node-level performance degradation. To mitigate the problem, they used *hugetlbfs*. They also employed *libhugetlbfs* [8], a wrapper library that semi-transparently maps application's text, data, and heap to a memory area backed by *hugetlbfs*. Their approach allowed Linux to achieve a performance comparable to CNK, both at the node level and systemwide. However, *hugetlbfs* does not eliminate the TLB

misses completely, so they can still be a performance problem for some applications. This approach also does not help with programming the DMA engine on BG/P. Moreover, the approach requires dynamic linking, while on Blue Gene almost all executables are statically linked, since this is the compiler default on that platform. The authors also found that dynamic linking introduced an overhead on accessing floating-point constants.

Navarro et al. [16] designed an effective transparent *superpage* management system that utilizes larger physical pages to reduce TLB misses and implemented it in FreeBSD on the Alpha processor. When a page fault occurs, the size of the superpage is chosen, and a set of contiguous page frames that covers the superpage is allocated from the buddy allocator. Fragmentation control and superpage promotion are part of their design. They evaluated their implementation with both benchmarks and realistic workloads and observed a 30% to 60% performance improvement. They targeted a relatively general-purpose usage, however, whereas we are focusing on high-performance applications on massive parallel machines.

3 Early Performance Evaluation

In this section, we present the results of a number of performance measurements we made on a mostly unmodified BG/P Linux kernel.

A popular misconception is that Linux would not be suitable as a compute node kernel because of a high level of operating system noise. We have disproved this notion on several older platforms [2, 4].

Figure 1 shows the results obtained by using our OS noise measurement benchmark Selfish [3] on BG/P. Essentially, the benchmark is a tight, busy loop that records anomalies in its execution time caused by OS interrupts (i.e., system noise). In this case, the Linux kernel with 4 KB pages and 100 Hz timer interrupt used only 0.027% of the CPU time, leaving 99.963% of the cycles to the benchmark. The interruptions are small and predictable, making them fairly easy to control.

Table 1 also shows that Linux OS noise is low. The FPU benchmark is basically a tight loop of FMA instructions. The benchmark kernel easily fits in the L1 instruction cache. On IBM’s CNK, the result matches the theoretical peak performance of BG/P floating-point unit per core ($4 \text{ flops} \times 850 \text{ MHz}$ yields 3.4 Gflops). Linux is only 0.09% slower. This is higher than the noise level measured earlier, probably because the noise measurement experiment does not account for noise events smaller than $1 \mu\text{s}$.




Figure 1: Noise on BG/P Linux

Table 1: CNK vs Linux: FPU benchmark

CNK [Gflops]	Linux [Gflops]	Performance Loss [%]
3.400	3.397	0.09

3.1 Memory Benchmarks

We evaluated the memory performance using two benchmarks: a streaming copy benchmark and a random memory access benchmark. Under Linux, we tested using two different page sizes: 4 KB and 64 KB.

The streaming copy benchmark is simple: it allocates a memory buffer, divides it in two, and copies data from one region to the other. The results for several different buffer sizes can be found in Table 2. With a 16 KB buffer size, which fits well in the L1 cache, Linux kernel incurs approx. 1.2% performance loss. This is higher than the overhead measured earlier; since no TLB misses occur during the execution of this benchmark, we suspect that Linux timer interrupts thrash the L1 cache. With a 256 KB buffer size, which is the maximum size that 4 KB TLB entries can cover ($64 \text{ TLB entries} \times 4 \text{ KB}$), we observe that the performance under Linux with 4 KB pages drops. We do not see any significant performance drop at 4 MB, even though this is the maximum size that 64 KB TLB entries can cover. This is because the overhead of a TLB miss becomes very small: updating a TLB entry takes approx. $0.3 \mu\text{s}$, while loading a 64 KB page takes approx. $60 \mu\text{s}$, which yields a 0.5% performance loss.

Table 3 presents the results of our random memory access benchmark. This benchmark reads data from a 16 KB, 256 KB, 4 MB, or 16 MB memory region in a random fashion. The results show that the performance

drops dramatically at 256 KB on Linux with 4 KB pages and at 4 MB on Linux with 64 KB pages. Many applications deal with buffers larger than 4 MB, so clearly, a more radical solution is needed than a mere increase of the page size by an order of magnitude.

Table 2: CNK vs Linux: Streaming copy benchmark

Size	CNK [MB/s] (σ)	Linux 4 KB [MB/s] (σ)	Linux 64 KB [MB/s] (σ)	Performance Loss	
				4 KB [%]	64 KB [%]
16 KB	2158.00 (0.03)	2131.67 (0.30)	2130.81 (0.40)	1.22	1.26
256 KB	1037.13 (0.03)	993.04 (6.84)	1031.44 (0.18)	4.25	0.55
4 MB	1037.39 (0.00)	993.51 (1.81)	1028.49 (2.50)	4.23	0.86
16 MB	1037.38 (0.02)	992.96 (1.73)	1028.08 (2.00)	4.28	0.90

Table 3: CNK vs Linux: Random memory access benchmark

Size	CNK [MB/s] (σ)	Linux 4 KB [MB/s] (σ)	Linux 64 KB [MB/s] (σ)	Performance Loss	
				4 KB [%]	64 KB [%]
16 KB	810.29 (0.00)	807.34 (0.23)	807.90 (0.08)	0.36	0.29
256 KB	810.36 (0.00)	67.16 (0.00)	810.16 (0.03)	91.71	0.02
4 MB	186.47 (0.00)	22.11 (0.13)	54.23 (0.02)	88.14	70.92
16 MB	62.60 (0.00)	17.76 (0.00)	22.66 (0.00)	71.63	63.80

3.2 NAS Parallel Benchmarks – Serial

To study the single node performance under different loads, we ran NAS Parallel Benchmarks (NPB) [1] version 3.3 in serial mode using the class A problem size. The executables were compiled using the IBM XL compiler and we could use the same binaries for both CNK and Linux. The results are in Table 4.

The IS result shows a large performance gap between CNK and Linux. This benchmark performs random memory accesses on a larger buffer, so the performance loss is close to the numbers from Table 3. Surprisingly, the CG benchmark on Linux with 64 KB pages runs faster than on CNK. We found that CNK enables the L2 optimistic prefetch, while Linux does not. In this particular case, the prefetch results in a performance loss. Brunheroto et al. [5] describe the L2 optimistic prefetch on Blue Gene/L and they confirm that the prefetch negatively affects the CG, IS, and SP benchmarks.

3.3 System Calls

Table 5 shows a comparison of the system call performance. This is one of the benchmarks where Linux outperforms CNK. The cost of a system call on CNK is almost eight times higher than on Linux. This is because CNK saves all possible registers on entry, while Linux saves a minimum set. Additional register saves and

Table 4: CNK vs Linux: NAS Parallel Benchmarks – serial

Type	CNK [Mop/s]	Linux 4 KB [Mop/s]	Linux 64 KB [Mop/s]	Performance Loss	
				4 KB [%]	64 KB [%]
BT	315.30	277.18	298.21	12.09	5.42
CG	37.15	36.21	37.88	2.53	−1.97
EP	3.23	3.03	3.18	6.19	1.55
FT	236.35	203.42	218.60	13.93	7.51
IS	23.51	5.47	5.86	76.74	75.07
LU	371.02	288.03	334.24	22.37	9.91
MG	254.54	238.45	250.20	6.32	1.71
SP	224.86	191.58	217.34	14.80	3.34

restores take time and potentially also thrash the L1 cache. A system call such as `gettimeofday` is often heavily used by applications or profiling software to manage timers, so it might impact the performance.

Table 5: CNK vs Linux: System call performance

	CNK [μ s/call]	Linux [μ s/call]	Performance Loss [%]
<code>gettimeofday</code>	3.91	0.51	− 86.96

4 Big Memory

We first briefly discuss standard paged memory management. We then explain our approach using “Big Memory.”

4.1 Standard Linux Memory Management

The Linux kernel is a virtual memory operating system. The main purpose of virtual memory is process isolation, but virtual memory also provides other optimizations or functionality such as copy-on-write, file caching, and memory swapping. Nowadays, virtual memory is considered mandatory in general-purpose operating systems.

Figure 2 provides an overview of memory management in Linux. The address space of each process consists of a set of virtual address ranges called virtual memory areas (VMAs). VMAs are created when a new process is started; the `mmap` system call can also create a new VMA. Creating a VMA is not equivalent to physical memory allocation; this takes place only on the first memory access within the VMA. Memory access attempts outside the VMAs result in a memory fault. VMAs have access permissions associated with them; incorrect access attempts result in memory faults as well.

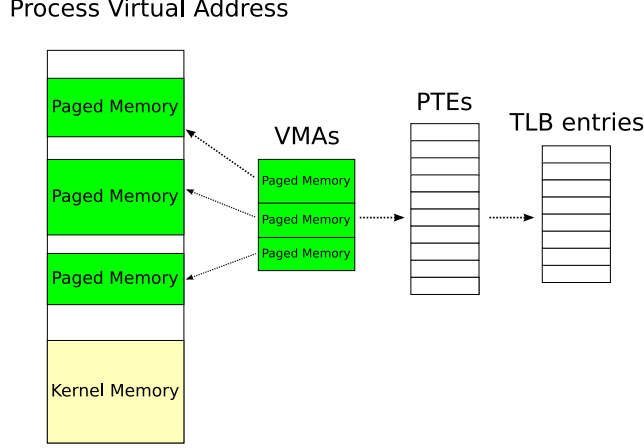


Figure 2: Memory Management in Linux

On most modern processors, user programs run inside virtual address spaces. In other words, user programs cannot address physical memory directly—the processor always has to convert a virtual address to the corresponding physical address. Page table entries (PTEs) are used for that purpose. The operating system kernel is responsible for creating PTEs based on the VMAs and for storing them in memory in advance or upon request. PTEs are stored in the kernel space (in main memory), so accessing a PTE is simply an overhead. To alleviate this overhead, modern processors cache recent address conversions in the translation lookaside buffer. The TLB definitely improves memory access speed, but it is not a complete solution because TLB entries are scarce. When a TLB miss takes place, a TLB entry needs to be loaded from a PTE.

The TLB entries associated with paged memory are flushed by the Linux kernel when the kernel switches to another process or when an associated VMA is removed. PowerPC Linux statically pins down a few TLB entries to cover kernel memory space, so TLB entries associated with kernel memory are not flushed when switching context.

4.2 Our Approach

Our basic idea is to provide applications with special memory regions that are covered by larger pages; we refer to these regions as *Big Memory*. Unlike the *hugetlbfs* regions discussed in Section 2, which use pages of intermediate size (2–4 MB) and only *reduce* the probability of TLB misses, Big Memory uses pages so large that the TLB misses are *eliminated* while the process is scheduled in.

We focus on the compute node environment here; we do not consider a general-purpose solution. The compute node environment has unique characteristics: a computational process tends to monopolize the

CPU execution unit, floating-point unit, memory, network, and other resources. In order to achieve peak performance, a single computational thread per core is preferable, since the kernel has to save the current context (all CPU registers associated with the current process, including double-floating-point registers in the case of the PowerPC 450), load new context, and invalidate caches at context switching time, thereby incurring substantial overhead, especially for CPU-intensive applications. Our approach is to allow only a single computational process to run at a time, to let the Linux kernel pin down TLB entries to cover the Big Memory area on the first access after the computational process has been scheduled in, and to remove them when that process gets scheduled out (to keep the memory mapping private to the computational process). Thus, an application does not suffer TLB misses when it accesses its Big Memory area.

Physical memory needed for the Big Memory area is reserved at boot time and is thus not available for use by the kernel as a regular, paged memory. It can be used only by a special computational process. This dramatically reduces the complexity of the implementation.

Unlike the `hugetlbfs`-based solutions, our approach is fully transparent, requires no code changes to the application, works with static executables, and covers all the necessary application segments.

The Linux kernel automatically prepares the Big Memory area for a computational application. How does the kernel determine which process should use Big Memory? Our solution is to alter the application's executable file; we use the `e_flags` field in the ELF header, which is reserved for processor-specific data. We defined a custom flag and wrote a tool that toggles it. We refer to the executables with the flag set as Zepto compute binaries, or ZCBs.

4.3 ELF Binary Interpreter

To load a ZCB into Big Memory, we have modified the Linux kernel ELF binary interpreter, specifically, the `load_elf_binary` function, which is invoked from the `execve` system call.

First, the ELF header is examined to see whether the binary being loaded is a ZCB. If it is, the kernel sets a bit in the `personality` field in the task structure so that other kernel functions can easily determine that the process is a ZCB by accessing the `current` variable.

Then the kernel creates a virtual memory range for Big Memory, using a simple offset mapping. We currently use 256 MB pages to cover the application memory; in the future we will improve the granularity of the Big Memory area by using a combination of different page sizes.

Once the Big Memory mapping is initialized, the kernel temporarily installs the Big Memory TLB

entries to copy the contents of both the command-line arguments and the environment variables to the application stack. The kernel also loads the application’s text and data sections to Big Memory instead of using the usual file mapping; the Big Memory mapping cannot be used for file mappings because it bypasses the Linux page allocator (see below). In other words, the entire big memory area is populated when the application binary is loaded.

4.4 Memory Manager for Big Memory

Our kernel reserves one VMA to cover Big Memory, and our internal memory manager takes care of memory chunks within the Big Memory area (heap, stack, etc.). To keep track of the memory chunks for `mmap` requests in Big Memory, the manager utilizes the kernel’s red-black tree—a structure normally used for managing entire VMAs. The red-black tree is a self-balancing search tree, which can be searched in $O(\log n)$ time, where n is the total number of elements in the tree.

The ZCB process address space is actually hybrid; it contains both regular paged memory and Big Memory (see Figure 3). Note that the Big Memory VMA does not have any associated PTEs, since the physical addresses of memory in that region are fixed. The behavior of the `mmap` system call varies depending on the request type. Anonymous, private requests, as used for large C library `malloc` calls, go to Big Memory and are tracked by our internal memory manager. On the other hand, file-backed mapping requests—used, for example, to support shared libraries—simply go to the regular Linux paged memory manager because Big Memory cannot be used for file mappings.

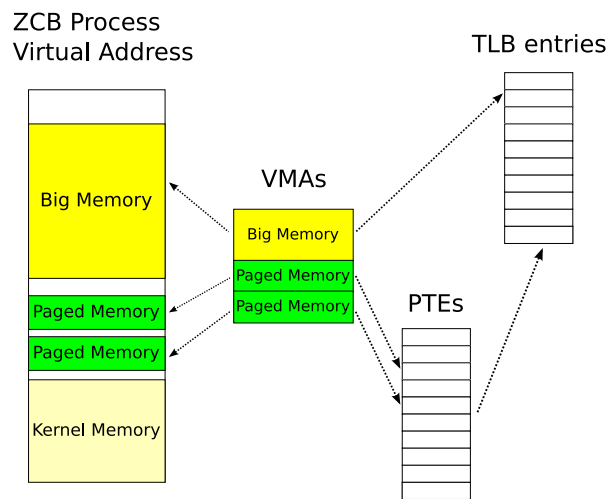


Figure 3: ZCB Process Address Space

4.5 Page Fault Handler

While we did not need to modify the TLB handler to implement Big Memory, we have added a hook to the Linux page fault handler (see Figure 4). The added code first checks whether the current task is a ZCB. If it is, the code checks whether the faulting address is within the Big Memory area; if so, it installs the Big Memory TLB entries. Essentially, we get a single TLB miss on the first access after the process has been scheduled in; the entries normally remain in place until the process is scheduled out again. With context switches being fairly rare on the compute nodes, the entries are semi-static.

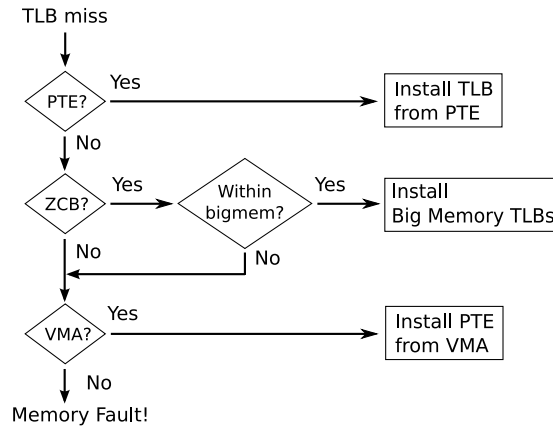


Figure 4: Big Memory Fault-Handling Flow

As shown in Figure 5, our Linux kernel partitions TLB entries in four groups: kernel mapping, paged memory, Big Memory, and device mappings. The last two groups are specific to our implementation. Currently, the number of TLB entries required by Big Memory is proportional to the area size; for example, seven entries are needed to cover 1792 MB. Three entries are used to cover the kernel memory (kernel low memory, to be precise). For efficiency, we also pin down some entries to cover BG/P-specific memory-mapped I/O devices: the collective network, the torus DMA, the lockbox, the Universal Performance Counter, and the Blue Gene Interrupt Controller. The Blue Gene Common Node Services code segment is also pinned.

4.6 Other Kernel Modifications

Introducing a new concept invariably causes issues; the Big Memory implementation is no exception. Big Memory breaks a number of assumptions in Linux memory management. A ZCB process is forked from a regular Linux process; however, it cannot call `fork` itself, since the `fork` code depends on the copy-on-




Figure 5: TLB Partitioning in the Zepto Kernel

write technique. The lack of `fork` may be acceptable when we focus on a high-performance compute node environment; besides, this is still better than the CNK, which does not support `fork` at all and must be rebooted to start another process.

Another issue is that the Big Memory mapping is strictly private and cannot be addressed from other processes' context. The Linux kernel function `access_process_vm` did not work for ZCB processes; this was an important issue because that function is used by the `ptrace` system call, on which debugging tools such as *strace* and *gdb* depend. Luckily, the problem was relatively easy to fix; basically, we now temporarily install the Big Memory mapping while the function is running, to allow other processes to access the address space of the ZCB process.

5 Performance Evaluation

To evaluate single-node performance on Linux with Big Memory, we ran our memory microbenchmarks and the serial implementation of the NAS Parallel Benchmarks. For parallel performance evaluation, we used (the parallel implementation of) NPB and the Parallel Ocean Program (POP). The experiments were performed on BG/P compute nodes with three different kernels: IBM CNK, Linux with 64 KB pages, and our enhanced Linux with Big Memory support. We used the IBM XL compiler that supports PowerPC 450 instructions to compile the benchmarks.

Sequential binaries are compatible between the kernels. Parallel binaries, however, could only be run on CNK and Linux with Big Memory because we do not have MPI support for paged memory. To make

the comparison as fair as possible, we rebuilt communication libraries for CNK and Linux from the same sources: BG/P-patched MPICH version 1.0.7 and Deep Computing Messaging Framework library (DCMF) version 1.0. Differences between the kernels result in slightly different binaries of the communication libraries and consequently in different parallel binaries, but the intermediate application object files were the same.

5.1 Memory Benchmarks

In Section 3.1, we showed the results of both the streaming copy benchmark and the random memory access benchmark; Linux had serious performance issues with random memory access.

We ran the memory benchmarks on Linux with Big Memory (“Bigmem”). The results are in Tables 6 and 7; they show that Big Memory completely eliminates the performance problem of random access to paged memory. Big Memory also slightly improves the streaming access pattern. Interesting, we observe that with 16 KB and 256 KB buffer sizes the performance loss with Big Memory is slightly larger than with 64 KB pages. We do not expect to see any significant numbers of TLB misses in either case; instead, we suspect that this is caused by the scheduler overhead. To keep Big Memory mapping private to the computational process, the associated TLB entries are flushed and re-installed when the process is scheduled in and out. The cost of removing Big Memory TLB entries is slightly higher than the cost of a single TLB miss. We also note that the performance loss with the streaming access is proportional to the memory bandwidth: 0.6% performance loss per GB/s.

Table 6: CNK vs Linux: Streaming copy benchmark

Size	CNK [MB/s] (σ)	Linux 64 KB [MB/s] (σ)	Linux Bigmem [MB/s] (σ)	Performance Loss	
				64 KB [%]	Bigmem [%]
16 KB	2158.00 (0.03)	2130.81 (0.40)	2130.15 (0.42)	1.26	1.29
256 KB	1037.13 (0.03)	1031.44 (0.18)	1030.85 (0.06)	0.55	0.61
4 MB	1037.39 (0.00)	1028.49 (2.50)	1031.67 (1.52)	0.86	0.55
16 MB	1037.38 (0.02)	1028.08 (2.00)	1031.87 (2.03)	0.90	0.53

Table 7: CNK vs Linux: Random memory access benchmark

Size	CNK [MB/s] (σ)	Linux 64 KB [MB/s] (σ)	Linux Bigmem [MB/s] (σ)	Performance Loss	
				64 KB [%]	Bigmem [%]
16 KB	810.29 (0.00)	807.90 (0.08)	807.04 (0.06)	0.29	0.40
256 KB	810.36 (0.00)	810.16 (0.03)	809.51 (0.47)	0.02	0.10
4 MB	186.47 (0.00)	54.23 (0.02)	186.18 (0.01)	70.92	0.16
16 MB	62.60 (0.00)	22.66 (0.00)	62.52 (0.00)	63.80	0.13

5.2 NAS Parallel Benchmarks – Serial

In Section 3.2, we showed the results of experiments with the serial implementation of NPB on both CNK and paged Linux.

Table 8 adds the results for Linux with Big Memory support. The results show that Big Memory clearly improves the memory performance, e.g., the huge loss observed with the IS benchmark is gone. On the other hand, the performance of the CG benchmark is now in line with the CNK, because unlike paged memory, Big Memory uses the same optimistic prefetch configuration as the CNK.

Table 8: CNK vs Linux: NAS Parallel Benchmarks – Serial

Type	CNK [Mop/s]	Linux 64 KB [Mop/s]	Linux Bigmem [Mop/s]	Performance Loss	
				64 KB [%]	Bigmem [%]
BT	315.30	298.21	314.61	5.42	0.22
CG	37.15	37.88	37.12	−1.97	0.08
EP	3.23	3.18	3.22	1.55	0.31
FT	236.35	218.60	236.22	7.51	0.06
IS	23.51	5.86	23.45	75.07	0.26
LU	371.02	334.24	370.67	9.91	0.09
MG	254.54	250.20	254.01	1.71	0.21
SP	224.86	217.34	225.01	3.34	−0.07

5.3 NAS Parallel Benchmarks – Parallel

So far we have shown that the Big Memory implementation definitely improves memory performance of applications on a single node. To evaluate parallel performance, we ran NPB in parallel mode on both CNK and Linux with Big Memory support.

We ran all the NPB benchmarks on both 1024 and 4096 nodes (except IS, which supports at most 1024 nodes), in SMP mode (one process per node), using class D problem size.

The results are shown in Table 9. The performance is very close: Linux runs were slower in all cases by 0.1–1.3%. The performance loss with the LU and CG benchmark increases at 4096 nodes, while the loss with the MG, FT, SP, and BT benchmarks decreases. We have not yet done a detailed investigation of these trends. We suspect that for benchmarks that improve as we scale up, the system noise might get hidden thanks to the hardware-assisted collective operations and DMA transfers (which are not affected by noise) or thanks to imperfectly balanced computations that force the majority of processes to wait doing nothing anyway. The performance loss with the EP benchmark stays pretty much independent of the number of nodes, since the benchmark is an embarrassingly parallel random-number generator. The 1.3% performance

loss observed is consistent with the loss we saw earlier with the streaming copy benchmark (see Table 6).

Table 9: CNK vs Linux: NAS Parallel Benchmarks – Parallel

Type	1024 Nodes			4096 Nodes		
	CNK [Mop/s/proc]	Linux Bigmem [Mop/s/proc]	Performance Loss [%]	CNK [Mop/s/proc]	Linux Bigmem [Mop/s/proc]	Performance Loss [%]
BT	304.03	300.48	1.17	241.53	240.81	0.30
CG	28.18	28.13	0.18	14.35	14.31	0.28
EP	3.78	3.73	1.32	3.78	3.73	1.32
FT	212.57	211.84	0.34	90.96	90.88	0.09
LU	288.99	287.84	0.40	211.99	210.62	0.65
MG	241.15	238.54	1.08	213.86	212.23	0.76
SP	149.24	148.58	0.44	132.60	132.34	0.20

5.4 Parallel Ocean Program

Parallel Ocean Program (POP) [12, 13] is a well known parallel application for studying the ocean climate system. We include it here as an example of a realistic real-world application.

We used unmodified POP version 2.0.1 from the POP website and the X1 benchmark data set. We ran it at the scale of 64 to 4096 nodes, in SMP mode. To be able to run using different node counts, we adjusted the `block_size_x` and `block_size_y` parameters, keeping both `max_blocks_clinic` and `max_blocks_tropic` at 1.

The results are shown in Table 10. Somewhat surprisingly, Linux runs scale better than CNK! We investigated this and found the earlier discussed high system call overhead of CNK to be responsible. POP calls `gettimeofday` about 100,000 times on each node, irrespective of the total number of nodes, resulting in a constant 0.36 second overhead. As the run times decrease with an increasing node count, this overhead becomes ever more significant, eventually allowing Linux to overtake CNK.

Table 10: CNK vs Linux: POP execution time

Nodes	CNK [s]	Linux Bigmem [s]	Performance Loss [%]
64	196.62	197.26	0.33
128	105.69	105.59	−0.09
256	57.37	57.00	−0.64
512	34.98	34.39	−1.69
1024	22.37	21.87	−2.24
2048	16.74	16.32	−2.51
4096	14.54	14.10	−3.03

5.5 LOFAR Online Central Processing

We had the opportunity to observe the performance of Big Memory in a real-life application. LOFAR is a radio telescope being built in the Netherlands [6]. In contrast to current radio telescopes that employ custom-built hardware as a correlator, LOFAR uses a Blue Gene/P supercomputer. Our previous work discusses the LOFAR central processor in more detail [17, 11].

LOFAR stations stream UDP/IP data directly into the Blue Gene/P I/O nodes at a rate of slightly more than 3 Gbps. These I/O nodes store the data in a main memory ring buffer, which is used to absorb network delays or temporary hiccups in the processing pipeline. From here the data is transported to the compute nodes, where the information is correlated.

Poor main memory performance of the I/O node running the default Linux kernel proved to be one of the major bottlenecks in trying to achieve optimum data throughput.

Table 11 shows a breakdown of the tasks on the Blue Gene/P I/O nodes running the LOFAR online processing application in a fairly standard 16-bit observation mode and a more challenging 4-bit mode. The required CPU resources, in processor cores, are shown for the stock IBM I/O node kernel in the 16-bit mode and for a ZeptoOS compute node kernel, modified to run on the I/O nodes, for both the 16-bit and 4-bit modes. We see that using the original I/O node kernel would require almost two entire I/O node cores just to handle the ring buffer, and that there are not enough compute resources available to perform all the tasks in real time, even for the less demanding 16-bit mode.

We used a slightly modified ZeptoOS compute node kernel, including support for the Ethernet device and excluding compute node specific devices like the torus network, on the I/O node. We reserved 1536 MB of main memory as the Big Memory area and used it for the ring buffer using six 256 MB TLB entries. The I/O node application was also adapted to copy 128 UDP/IP packets into the ring buffer at once, instead of one at a time. These two optimizations reduce the resources required to copy data into the ring buffer by more than 600%, about half of which can be attributed to the lack of TLB misses in the Big Memory area.

The more challenging 4-bit mode significantly increases the potential number of TLB misses. Although limitations in the default kernel make direct comparison impossible, smallscale memory benchmarks show a more than 500% reduction in required CPU cycles using Big Memory when copying UDP/IP packets one by one.

Clearly, the access pattern of the LOFAR I/O node application is very susceptible to performance hits caused by TLB misses. With the stock I/O node kernel, the processor was unable to achieve our throughput

requirements. Preventing TLB misses for at least part of the application, combined with several other optimizations not discussed here, allowed us to reduce CPU load considerably, increasing I/O node performance to well beyond our original requirements.

Table 11: LOFAR I/O node processing (#cores)

	Stock	ZeptoOS	
	16 bit	16 bit	4 bit
Receive UDP/IP packets	1.44	1.44	1.22
Copy data to ring buffer	1.80	0.27	0.37
Send ring buffer to CN	1.40	0.52	0.61
Receive data from CN	1.00	0.10	0.35
Send results to storage	0.40	0.32	0.64
Total system load	151%	66.5%	79.7%

6 Conclusions

This paper presented the implementation of Big Memory support for BG/P Linux—a transparent, flat memory space for computational processes. Big Memory addresses two major issues encountered when attempting to run high-performance code on the BG/P Linux: poor memory performance caused by TLB misses handled in software, and the difficulties of writing an efficient communication stack caused by the limitations of the BG/P torus’ DMA engine.

Our OS noise measurement benchmark shows that Linux kernel with 4 KB pages and 100 Hz uses only 0.027% of the CPU time. The FPU benchmark, which is a tight loop of FMA instruction, suffers 0.09% of the performance loss compared to CNK. We think that the 0.09% loss is due to the time spent executing the OS tick interrupt handler.

Our experiments have shown that memory benchmarks running under a standard Linux kernel with 4 KB pages suffer up to 92% of performance loss compared to CNK. Increasing the Linux kernel page size to 64 KB narrows the gap; the worst case performance loss drops to 71%. Once we introduce Big Memory support, the worst case performance loss drops to just 1.3%.

We ran NAS Parallel Benchmarks at up to 4096 nodes and confirmed that Linux has no problems at this scale. In fact, with the majority of the benchmarks, the performance loss decreases with an increasing total number of nodes. We also ran the Parallel Ocean Program (POP) and the results show that for this application Linux scales better than does CNK. We found that the high overhead of system calls under CNK is responsible.

Employing Big Memory on the I/O nodes was instrumental in reducing the I/O node CPU resources required for LOFAR online central processing. A 500–600% performance increase was observed in key parts of the application, allowing the I/O nodes to achieve their required throughput.

Our modifications to the Linux kernel are relatively small, principally because we focused exclusively on the requirements of computational processes, rather than trying to solve the problem in a generally applicable way, which would have been far more complicated. We maintain two versions of Linux kernel, and we found porting the Big Memory patches between the kernels to be straightforward.

Contemporary parallel architectures such as Blue Gene/P, Cray XT5, and Roadrunner use commodity CPUs such as Intel Xeon, AMD Opteron, or IBM PowerPC, instead of designs dedicated to computational environments. The memory management units in these processors are essentially designed to support a highly multitasking environment. It would be interesting if future designs had hardware support for computational process address space similar to the Big Memory area that we implemented, to allow for a seamless coexistence of high-performance applications and standard Unix processes on the compute nodes.

Along with the Big Memory implementation, we conceived the idea of a special process that the kernel treats differently from other processes. In the case of Big Memory, the kernel creates a different application address space, and we showed that this idea works for compute nodes. We have also experimented with other uses of this feature, such as disabling nonessential interrupts in order to reduce the system noise when a computational process gets scheduled in.

Our current implementation is suitable for benchmarking and simple applications. We need to make several improvements in the quality of implementation, such as the granularity of the Big Memory area. So far, we have not observed any parallel performance degradation caused by OS noise at the scale of up to 4096 nodes. Confirming this at very large scales and investigating why OS noise does not seem to affect the HPC applications will be part of our future research.

Acknowledgments

We thank IBM’s Todd Inglett, Thomas Musta, Thomas Gooding, George Almási, Sameer Kumar, Michael Blocksome, and Robert Wisniewski for their advice on programming the Blue Gene hardware. We also thank our past summer interns Peter Boonstoppel, Hajime Fujita, Satya Popuri, and Taku Shimosawa, who contributed to the ZeptoOS kernel. Additionally, we thank Astron’s John W. Romein, who evaluated Big

Memory on the I/O nodes.

This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory.

References

- [1] D. Bailey et al. The NAS parallel benchmarks. *International Journal of High Performance Computing Applications*, 5(3):63–73, 1991.
- [2] P. Beckman, K. Iskra, K. Yoshii, and S. Coghlan. The influence of operating systems on the performance of collective operations at extreme scale. In *Proceedings of the 8th IEEE International Conference on Cluster Computing*, Barcelona, Spain, Sept. 2006.
- [3] P. Beckman, K. Iskra, K. Yoshii, and S. Coghlan. Operating system issues for petascale systems. *ACM SIGOPS Operating Systems Review*, 40(2):29–33, Apr. 2006.
- [4] P. Beckman, K. Iskra, K. Yoshii, S. Coghlan, and A. Nataraj. Benchmarking the effects of operating system interference on extreme-scale parallel machines. *Cluster Computing*, 11(1):3–16, Mar. 2008.
- [5] J. R. Brunheroto, V. Salapura, F. F. Redigolo, D. Hoenicke, and A. Gara. Data cache prefetching design space exploration for Blue Gene/L supercomputer. In *Proceedings of the 17th IEEE International Symposium on Computer Architecture on High Performance Computing*, pages 201–208, Rio de Janeiro, Brazil, Oct. 2005.
- [6] H. R. Butcher. LOFAR: First of a new generation of radio telescopes. In *Proceedings of SPIE*, volume 548, pages 537–544, Oct. 2004.
- [7] K. Chen, R. Seth, and H. Nueckel. Improving enterprise database performance on Intel Itanium architecture. In *Proceedings of the Linux Symposium*, pages 98–108, Ottawa, ON, Canada, July 2003.
- [8] D. Gibson and A. Litke. libhugetlbfs. <http://sourceforge.net/projects/libhugetlbfs>.
- [9] IBM Blue Gene. <http://www.research.ibm.com/bluegene/>.
- [10] IBM Blue Gene team. Overview of the IBM Blue Gene/P project. *IBM Journal of Research and Development*, 52(1/2):199–220, 2008.

- [11] K. Iskra, J. W. Romein, K. Yoshii, and P. Beckman. ZOID: I/O-forwarding infrastructure for petascale architectures. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 153–162, Salt Lake City, UT, Feb. 2008.
- [12] P. W. Jones, P. H. Worley, Y. Yoshida, J. B. White, III, and J. Levesque. Practical performance portability in the Parallel Ocean Program (POP): Research articles. *Concurrency and Computation: Practice and Experience*, 17(10):1317–1327, 2005.
- [13] D. J. Kerbyson and P. W. Jones. A performance model of the Parallel Ocean Program. *International Journal of High Performance Computing Applications*, 19(3):261–276, 2005.
- [14] J. E. Moreira et al. Blue Gene/L programming and operating environment. *IBM Journal of Research and Development*, 49(2/3):367–376, Mar. 2005.
- [15] J. E. Moreira et al. Designing a highly-scalable operating system: The Blue Gene/L story. In *Proceedings of the ACM/IEEE Conference on Supercomputing*, Tampa, FL, Nov. 2006.
- [16] J. Navarro, S. Iyer, P. Druschel, and A. Cox. Practical, transparent operating system support for superpages. In *Proceedings of the 5th ACM Symposium on Operating System Design and Implementation*, volume 36 of *ACM SIGOPS Operating Systems Review*, pages 89–104, Boston, MA, Dec. 2002.
- [17] J. W. Romein, P. C. Broekema, E. van Meijeren, K. van der Schaaf, and W. H. Zwart. Astronomical real-time streaming signal processing on a Blue Gene/L supercomputer. In *Proceedings of the 18th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 59–66, Cambridge, MA, July 2006.
- [18] E. Shmueli, G. Almási, J. Brunheroto, J. Castaños, G. Dózsa, S. Kumar, and D. Lieber. Evaluating the effect of replacing CNK with Linux on the compute-nodes of Blue Gene/L. In *Proceedings of the 22nd ACM International Conference on Supercomputing*, pages 165–174, Kos, Greece, July 2008.
- [19] ZeptoOS project. <http://www.zeptoos.org/>.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory (“Argonne”). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.